

NASA Technical Memorandum 107685

Current State and Future Direction of Computer Systems at NASA Langley Research Center

Computer Systems Technical Committee

September 1992

(NASA-TM-107685) CURRENT STATE AND
FUTURE DIRECTION OF COMPUTER
SYSTEMS AT NASA LANGLEY RESEARCH
CENTER (NASA) 34 p

N93-12278

Unclass

G3/61 0127233



National Aeronautics and
Space Administration
Langley Research Center
Hampton, Virginia 23665-5225

CONTENTS

Introduction	1
Real-Time Computer Systems	2
Embedded Systems	4
High Performance Computing.....	7
Distributed Computing Networks.....	10
Artificial Intelligence	12
Data Acquisition Systems.....	14
Visualization	22
Summary	30

INTRODUCTION

(Jerry H. Tucker and James L. Rogers)

Computer systems have advanced at a rate unmatched by any other area of technology. As performance has dramatically increased there has been a no less dramatic reduction in cost. This constant cost performance improvement has precipitated the pervasiveness of computer systems into virtually all areas of technology. This improvement is due primarily to advances in microelectronics. These advances have resulted in the capability to place millions of transistors on a single integrated circuit. This ability to build inexpensive and extremely complex integrated circuits has resulted in microprocessors that incorporate architectural features that only a few years ago were reserved for supercomputers. We can expect microprocessor performance to double about every two years for the remainder of the decade and by the end of the decade it should be possible to fabricate multi-microprocessor integrated circuits containing fifty to one hundred million transistors. Most people are now convinced that the new generation of supercomputers will be built using a large number (possibly thousands) of high performance microprocessors. Although the spectacular improvements in computer systems have come about because of hardware advances there has been a steady improvement in software techniques. We now have understanding of why programs fail and how to make them more reliable and maintainable. This insight has been translated into new languages and is effecting the evolution of software development techniques.

This paper discusses selected important areas of computer systems at NASA Langley Research Center. These include real-time systems, embedded systems, high performance computing, distributed computing networks, data acquisition systems, artificial intelligence, and visualization. Langley Research Center has been involved in real time flight simulation research for over fifty years, and is currently implementing the Flight Simulation Computing System (FSCS). The significant performance increase of the FSCS can be used to simulate systems that have not been previously simulated. Embedded microprocessor systems are used in a wide variety of flight and ground applications at Langley. The capability and sophistication of these system have increased dramatically reflecting the rapid advances in microprocessor technology. One application of embedded systems, discussed in this paper, is in custom wind tunnel automation systems. High performance

computing has been associated with CRAY-like supercomputers consisting of a small number of very high performance processors. These conventional supercomputers are, for some applications, now being challenged by massively parallel computers containing a large number of inexpensive microprocessors. This paper discusses distributed computer networks that allow users to connect to a wide variety of computers at NASA Langley as well as outside sources. Another area which has undergone change, is that of data acquisition systems. To help understand current state and future direction, the history of these systems is briefly described. A typical current data system is based on a monolithic design developed in the 1960's and implemented on a proprietary process control computer. The proposed future data system will utilize a new functional specification and will be based on a cooperating network of open architecture computers. Two areas discussed in this paper that are on the leading edge of software technology are artificial intelligence and visualization. The two main areas of research with artificial intelligence techniques at NASA Langley are knowledge-based systems and neural networks. Scientific visualization is important to many research disciplines at Langley, and is primarily supported by the Data Visualization and Animation Lab in ACD.

REAL-TIME COMPUTER SYSTEMS

(Billy R. Ashworth and Jeff I. Cleveland II)

Real-Time Computer Systems at LaRC include a variety of applications where computed results are required in "real-time". In terms of piloted flight simulation, the results of a control movement of a simulated vehicle must be available to the pilot/controller as soon as they would be in the "real world" (i.e., the movement of the simulated vehicle must be computed and presented to the pilot in the same time as the "real-time" vehicle would respond). The use of the real-time systems covers numerous applications at LaRC ranging from piloted real-time simulation to embedded flight software to data acquisition systems. This year's summary report will concentrate on the state of the technology in piloted real-time flight simulation and real-time data acquisition systems.

Real-time flight simulation utilizes computers connected to piloted cockpits to simulate aircraft or spacecraft flight. Three major subsystems comprising the simulation facility at LaRC are the

simulation computers with the real-time operating system, the real-time input/output system, and the attached real-time simulation sites consisting of cockpits, graphical display generating devices, and control consoles.

LaRC has been using real-time flight simulation to support aerodynamic, space, and hardware-in-the-loop research for over 50 years. In the mid-1960's LaRC pioneered the first practical real-time digital flight simulation system using large, high-performance general purpose computers. Initially, Control Data Corporation (CDC) 6600 computers were installed. In 1976, they were replaced with CDC CYBER 175 computers. In 1987, the analog-based simulation input/output system was replaced with a high-performance fiber optic-based digital network. With the increased complexity and higher performance requirements for the simulation of modern aircraft, LaRC is currently replacing the CDC computers with CONVEX Computer Corporation super computers.

The current trend in real-time computing involves the use of multiple processors; however, the power of these parallel processors is vastly different. Many training simulators and other less frequently changing simulations are beginning to use a distributed approach where basic components, such as engines, are simulated in separate microprocessors with the total simulation consisting of several interconnected microprocessors. The large, frequently changing simulations used at LaRC, use large central computers with powerful (multiple) processors. This choice is desirable where frequent changes to the program code are required for flight research. Having decided to continue using powerful, multi-CPU, centralized computers, LaRC issued a Request for Proposals in May 1989 and subsequently awarded a contract to CONVEX Computer Corporation in December of that year for the Flight Simulation Computing System (FSCS).

LaRC is currently implementing the FSCS. When completed, it will consist of one CONVEX C3850 computer containing five central processing units and one CONVEX C3230 containing three central processing units. The C3850 will support up to four independent simulations operating concurrently and the C3230 will support up to two simulations operating concurrently. A simulation may use as many real-time central processing units as there are available on a computer. Each of the CPU's is 3 to 5 times more powerful than the single CPU CYBER mainframes currently in use.

With the large increase in performance, the FSCS can simulate systems that have not been previously simulated. Advanced, highly-complex fighters with flexible aerodynamic structures can now be simulated with high fidelity. Complex weather models including the simulation of microbursts are being added to flight simulations including full mission air traffic control models. The increase in computer power provides the capability of simulating systems with higher solution frequency content. This includes the ability to simulate a control system for actively controlling space structures at rates of up to 1,000 cycles per second.

EMBEDDED SYSTEMS

(Jerry H. Tucker and Charles E. Niles)

The invention of the microprocessor in the early 1970's made it possible to build digital systems containing dedicated computers. These first microprocessors were slow, expensive, and cumbersome to use. Because of this, microprocessors were initially embedded in systems only if they could replace a substantial amount of conventional hard wired logic, and if speed was not a concern. The second generation of microprocessors, which were faster and simpler to use, arrived in the mid-seventies. It was at this time that microprocessors began to appear in a multitude of diverse applications. Today, the microprocessors used in embedded systems range from a simple single chip micro controller to high performance parallel processors.

Embedded microprocessor systems contain one or more dedicated microprocessors and typically differ from a conventional general purpose computer system in several respects: embedded systems usually do not contain computer peripherals; the embedded computer is dedicated to a single task (it continuously executes a single program which generally resides in read only memory); the interaction of the computer with a human may be nonexistent or very limited, and the existence of a computer in the system may not be obvious; most embedded systems are used in control applications which have limited computational requirements.

In the 1970's, the software for most embedded microprocessor systems was written in assembly language. In recent years the trend has been toward high level languages with little or no use of

assembly language. Older high level languages such as FORTRAN and PL/M have been largely replaced by C and to a lesser extent Ada. Despite having been developed for real time embedded systems, ADA's acceptance has been low. In the DOD, where its use is mandated, Ada is used in less than 50% of the embedded systems, and in the commercial sector its use is less than 7% for embedded systems.

The use of embedded microprocessor systems for space flight pose special problems. These problems include restrictions on size, power, and weight, as well as requirements for long term high reliability in harsh environments. One particularly troublesome concern for space flight embedded systems is the effect of radiation. This is because the rapid improvements in processor performance have been obtained primarily by the addition of architectural features which require a higher component density. This density is achieved by placing millions of extremely small components on a single integrated circuit. As component size decreases, the integrated circuit tends to become more vulnerable to radiation effects. Since state-of-the art development is driven by commercial demands there is only limited motivation for the leading microprocessor manufactures to address the radiation problem. The effect is that the space flight system designer is forced to use technology that lags several generations behind state-of-the art. A technological development that should improve this situation is application specific integrated circuits (ASIC's). With ASIC technology a designer uses software to describe, synthesize, and simulate a digital system. This system can then be implemented in a single integrated circuit. Replacing an entire system or a large portion of a system with a single ASIC will result in reduced weight and size, as well as reduce power consumption. Because there are fewer elements to fail, and radiation effects can be minimized by implementing the ASIC using a radiation hard process, reliability should improve. It is expected that by the year 2000 it will be possible to fabricate ASIC's containing 40 million transistors. With this capability, ASIC technology should provide NASA with state-of-the art capability for both flight and ground based applications.

Facility Automation

Embedded systems used in custom wind tunnel automation systems are becoming more generalized as technology improves. In the 1970's, these systems consisted of a set of boards and firmware that interfaced to hardware control panels and device actuators. The

firmware included a custom operating system, device drivers, and control algorithms all rolled into one application program. In the 1980's, the set of boards became faster and the application in firmware became larger as hardware control panels were replaced by software-oriented operator input and display devices. Since 1985, real-time multi-tasking operating systems have removed mundane functions from application and opened the doors to reusable software components. Software architecture has become cleaner with more clearly defined interfaces between tasks.

Systems to be deployed within the next year at LaRC offer many improvements over their predecessors. The baseline target platform relies on the Multibus II architecture, an 8086 family CPU board, intelligent input/output boards, and the RMX-III real-time, multi-tasking operating system kernel. A baseline set of application tasks executing under RMX-III provides a layer known as a "multi-tasking control environment" through which facility-specific control algorithms obtain control information (process data, set points, and control parameters) and issue actuation commands. Modular design allows systems to interface with existing hardware control panels or touch screen devices mounted on high-resolution color monitors. Most systems support control functions from a supervisory computer via an ETHERNET or serial link.

Although RMX-III has been an asset to embedded systems development, it has been an inconvenience due to the lack of quality development tools. Consequently, future systems will be based on a different operating system which is UNIX-like, complies with POSIX standards, runs on multiple hardware platforms, and provides virtually the same development and target environments. Future operator interfaces will be X-Window applications, where clients reside on the embedded system. In complex systems, where operator interface and control functions are distributed among a workstation, a supervisory controller, and two or more embedded controllers, the operator interface will connect to a local network to interface to control functions.

HIGH PERFORMANCE COMPUTING

(Jerry H. Tucker and Geoffrey M. Tennille)

The classification of a computer as a "supercomputer" is a subjective judgment. It comes not from specific, well-defined quantitative measurements of a computer's speed but from its power relative to the other high-speed computers of its day. If the scientific community accepts that a particular computer can solve the largest of problems that are too complex for the vast majority of other scientific computers available, then it is classified as a supercomputer. Consequently, any quantitative measure used today to define supercomputer speed would be inaccurate in several years.

Historically, supercomputers have been characterized by large fast-access memories, very fast computational rates, usually achieved by innovative architectural features, and high cost. At LaRC we have a 128 million word (Mwd) CRAY-2, called Voyager, and a 128Mwd CRAY Y-MP, called Sabre. Both of these computers have multiple Central Processing Units (CPU's), each CPU being itself a very fast vector processor. Voyager has 4 CPU's and Sabre has 5 CPU's. Vector processing, while common on today's high speed computers, was the distinguishing feature back in the 1975 era when it provided the extra power to establish the CDC CYBER-205 and CRAY-1 computers as supercomputers. Vector processing was introduced when it became clear that the current path, increasing speed through faster electronics, would not provide the improvement needed to solve the next generation of problems. Vector processing did not make any one computation faster, but used the architectural trick of overlapping the computation of a large number of operands to reduce the average time per result.

The next innovative architectural step, taken by Cray Research, was to introduce multiple processors that could either provide multiplicative gross power, or be used together to solve a single problem at this higher gross rate. This was the first practical use of parallel processing, albeit on a very coarse scale. However, this high-level parallelism set the stage for the next advance in computer technology, massively parallel computing. Aside from the introduction of coarse parallel processing, most of the other advances in supercomputing rates over the last 10 years have resulted from increased clock speed, increased memory size and access-speed, the use of multiple functional units within a single CPU (that is, a capability to do a vector add and a vector multiply

simultaneously), and an improvement in compiler technology to automatically recognize vector instructions. Additionally, there has been a modest increase in the number of processors (up to 8 for the CRAY Y-MP) as well as improvements in the implementation of parallel processing. Each CRAY Y-MP CPU is capable of 320 Million Floating Point Operations per second (Mflop/s) so that a fully configured machine (8 CPU's) has a peak performance of 2500 Mflop/s (denoted 2.5 Gflop/s).

LaRC is one of three NASA centers (Ames and Lewis are the others) participating in the Computational AeroSciences (CAS) Grand Challenge of the Federal High Performance Computing and Communications Program (HPCCP). The thrust of CAS is to develop interdisciplinary applications for the design of advanced aerospace vehicles. The CAS project is expected to place an additional Massively Parallel Processor (MPP) testbed at LaRC. This testbed will permit experimentation with both heterogeneous distributed computing among local computers and long distance cooperation among tasks executing on a similar architecture at a remote site, such as Lewis Research Center. The NASA CAS centers are also participants in the Computer Science Corporation consortium for the Intel Touchstone Delta system at California Institute of Technology.

We have reached a point where it appears that the current CRAY-like model for supercomputing will provide only modest increases for the future. For instance, the next generation CRAY computer, the CRAY C-90, will offer a factor of three performance improvement per CPU over the CRAY Y-MP and an increase to 16 CPU's for a peak capability of 16 Gflop/s. For researchers trying to solve the next generation of grand challenge applications, this falls well short of the Teraflop/s rate (1,000 Gflop/s) that is needed. It is now widely believed that truly significant advances, at affordable prices, can be realized only through the introduction of thousands of moderately fast, relatively inexpensive microprocessors. This shift to "massively parallel processors (MPP)" has already started as evidenced by the increased number of commercial products now available (even Cray Research is designing a MPP system).

Within the scientific community, there are predictions of steadily increasing microprocessor performance. Numerous university and government research programs point toward massively parallel systems and the software that will be necessary to effectively utilize them. Some government laboratories have even made bold

predictions that they do not ever expect to buy another CRAY-like supercomputer. Conventional CRAY-like, supercomputing is well understood today, there needs to be continued improvement in software utilities for debugging, performance analysis, and parallelization. However, needing far more attention, is the MPP world where the issues of programming and performance are of great importance.

Massively parallel computers may contain hundreds to thousands of processing elements, none of which is a supercomputer itself. At the current rate of increase in the theoretical speed of these so-called "killer-micros", their peak performance is expected to exceed that of current conventional supercomputer CPU's in the near future. Thus by the middle to late 1990's it should be possible to build a MPP computer, with a peak Teraflop/s speed, using only a few hundred or thousand processors. Current experience however, indicates that peak speeds of 5 to 10 Teraflop/s may be necessary to sustain a one Teraflop/s performance. The low ratio of sustained to peak speeds on parallel processors is due in large part to immature compilers and communications techniques.

The transition to MPP computing is not expected to be easy. There are nearly as many ways to connect hundreds of processors together as there are vendors supplying this type of architecture. Examples of these interconnections are meshes, toroids, hypercubes and rings. Some MPP machines rely on a message-passing paradigm for programming, while others use a data-parallel method. Most MPP systems have distributed memories, but a few support shared memories. This wide diversity in architecture and style of programming presents both a stiff challenge as well as the potential for exciting software research. A number of commercial and research software packages exist or are being developed to assist in the porting of applications to MPP architectures. Their primary purpose is to provide a uniform programming environment across numerous platforms. A corollary development to this effort is heterogeneous computing, where several different machines are used simultaneously in the solution of a single problem. This requires synchronization between machines as well as run-time conversion and exchange of potentially non-compatible binary data.

Researchers at LaRC's ICASE have been investigating parallel processing for several years, beginning with the Flex-32 and continuing with the Intel iPSC/2 and the current 32-node iPSC/860.

Researchers at LaRC also have access to a 128-node iPSC/860 and a 32K-node CM-2, both at NASA Ames Research Center. Extensive work has been done in the development of message-passing primitives to facilitate the distribution of data for applications using unstructured grids and in compiler technology to support the data structures required for these applications. Much work still needs to be done in the area of software tools, especially in compiler technology for automatic exploitation of parallelism that is inherent within an application. Additionally, work is underway at LaRC to evaluate commercially available tools that provide assistance in the development and porting of applications.

To support the development of production level MPP computing, a substantial effort to develop benchmarks for parallel systems is underway at the Center. One problem, yet to be resolved, is how to translate the performance results obtained into evaluation criteria for machines that have not yet been built. Also, over twenty vendors have made presentations at LaRC to define their view of the current state-of-the-art in parallel processing. The management of the Central Scientific Computing Complex is reviewing the state of parallel processing to define both a strategy for bringing MPP computing into a production environment as well as to define the hardware and software requirements for such a system. Early experience with HPCCP testbeds is expected to provide substantial hands-on experience with the management and effective utilization of MPP systems.

DISTRIBUTED COMPUTING NETWORKS **(Charles E. Niles and John W. McManus)**

The proliferation of personal and small group computing resources at LaRC in the 1980's was brought about by the need for a more modern and diverse range of applications which were easier to use and responded more quickly. Since many users still required access to central or outside computing resources, they found ways to connect to (or around) the forerunner of today's LaRCNET.

LaRCNET presently provides the backbone which supports center-wide data communications and access to outside networks. This backbone includes physical media, bridges, gateways and software communications protocols (e.g. TCP/IP, DECNET). The network enables users to freely connect to mainframe and high performance

computers from IBM-compatible microcomputers, Unix workstations, Macintosh microcomputers, and other computers. Many users are well acquainted with remote terminal sessions via TELNET, file transfers using FTP, and transparent file storage via NFS.

Recognizing that network loading is an increasing problem, the Analysis and Computation Division (ACD) has presented plans to upgrade the backbone architecture including "fiber optic media to the desktop" and sub-networks in order to improve speed, increase throughput, and reduce security risks from outside influences.

ACD has also announced plans to establish a standardized E-mail service. The plans address provisions for a central post office (i.e. a default mailbox for each center employee) with essential mail functions (send, forward, etc.). Many technical issues of interoperability among different mail protocols operating on different platforms will be handled by the post office.

It appears that connectivity problems are under control and interoperability problems are slowly disappearing as commercial products gradually appear for different computing environments. Growing support for windowing systems, especially X-Windows, allows researchers and engineers to sit at their desktop computing resource and use remote computing resources transparently.

However, there exists a large void in distributed computing capabilities at LaRC dealing with real-time data access. Many facilities have antiquated, proprietary data collection systems which make it cumbersome to obtain data for on-line data analysis to enable researchers to visualize test activities. Also, these old systems supply data so slowly to automation systems that even rudimentary control is difficult, let alone advanced control. Replacing these old systems with POSIX-compliant systems which provide TCP/IP and X-Windows support would help fill the void. As an added benefit, these systems would look and act more like the researcher's desktop environment.

Some organizations have distributed computing needs which can only be achieved when financial and political issues are resolved. Other organizations have solved their immediate distributed computing needs. Their system administrators have knowledge and experience which should be tapped when organizations in need of technical expertise embark on the distributed computing path.

ARTIFICIAL INTELLIGENCE

(James L. Rogers and John W. McManus)

Currently, there are two main areas of research with artificial intelligence techniques at NASA LaRC, knowledge-based systems (KBS) and neural networks. The research in these two areas are performed to two types of hardware, (1) an AI specific workstation such as the Symbolics; or (2) a general purpose workstation such as the Sun or Iris, or a personal computer such as the IBM PC or Apple Macintosh. The languages supporting this research include Fortran, C, and Lisp.

Knowledge-Based Systems

The Human/Automation Integration Branch has developed a KBS for case-based reasoning about physical systems (Feyock, Karamouzis, and Schutte) such as diagnosing and handling in-flight malfunctions. This system examines a data base of problems to see if there is a precedent for the current problem and uses that information for the purposes of diagnosis, prognosis, and recovery planning. The system is written in Lisp and executes on the Symbolics, Iris, and IBM PC (486). Most of the work is migrating to the Iris and IBM with very little work being done on the Symbolics.

The Human/Automation Integration Branch is also performing research in aiding the flight crew of commercial transport aircraft (Schutte and Smith). Their approach is to provide human centered support for the flight crew rather than a technology-centered replacement for the flight crew. Their research is focused in two areas, fault management and information management. They have modified standard AI techniques to execute in real time to provide the crew with information which is useful, understandable, timely, and accessible. Their current software executes on Symbolics Lisp Machines. Their future hardware plans are to phase out the Symbolics and move to more conventional platforms such as IRIS, SUN, MAC, and IBM PC.

Vigyan and Old Dominion (Mehrotra and Wild) are investigating issues in the validation and verification of knowledge-based systems. This investigation is currently focused towards studying the factors critical to the grouping of rule bases, clustering analysis. Their code is in C and is executed on a Sun workstation.

The Interdisciplinary Research Office has developed a KBS for decomposing a complex design problem (Rogers). This tool orders a sequence of coupled subsystems by minimizing the amount of feedback among the subsystems, identifies interactions among the subsystems, and groups subsystems belonging to an iterative process. The data is displayed as a design structure matrix in NxN format. The knowledge base accesses the CLIPS system developed by NASA JSC. The system is written in C and executes on an Apple Mac and Sun Workstation.

The Aircraft Guidance and Control Branch is researching ways to apply mature programming techniques from the AI sub-field of computer science to aircraft guidance and control problems (McManus). AGCB is developing methods for integrating symbolic and numeric computation techniques for real-time applications. This research requires the development of real-time Knowledge-Based Systems (KBS), and an architecture for integrating them into conventional guidance and control systems. The branch has selected the blackboard system problem solving model and developed a set of design and analysis techniques for concurrent blackboard systems. Blackboard systems are a natural progression of Artificial Intelligence based systems into a more powerful problem solving technique. They provide a way for several highly specialized knowledge sources to cooperate to solve large, complex problems. Blackboard systems incorporate the concepts developed by rule-based and expert systems programmers and include the ability to add conventionally coded knowledge sources. The small and specialized knowledge sources are easier to develop and test, and are hosted on hardware specifically suited to the task that they are solving.

One blackboard system is called Paladin, a real-time Tactical Decision Generator (TDG) for air combat engagements. Paladin uses specialized knowledge-based systems and other AI programming techniques to address the modern air combat environment and agile aircraft in a clear and concise manner. Paladin is designed to provide insight into both the tactical benefits and the costs of enhanced agility.

The system was developed using the Lisp programming language on a specialized AI workstation. Paladin utilizes a set of air combat rules, an active throttle controller, and a situation assessment module that have been implemented as a set of highly specialized knowledge-based systems. The situation assessment module was

developed to determine the tactical mode of operation (aggressive, defensive, neutral, evasive, or disengagement) used by Paladin at each decision point in the air combat engagement. Paladin uses the situation assessment module and the situationally dependent modes of operation to more accurately represent the complex decision-making process of human pilots. This allows Paladin to adapt its tactics to the current situation and improves system performance.

Neural Networks

The Automation Technology Branch and Systems Architecture Branch have developed a system for extracting an input-output model of scalar dynamical systems from a trained neural network and have extended it to analytical nonlinear dynamical systems (Soloway and Haley). Preliminary results look promising. They plan to design a nonlinear control system using a real time neural network. Current operational code is written in C and runs on an IBM PC or Sun.

The Interdisciplinary Research Office has developed a system which couples a neural network (NETS) with an optimization program (Rogers and LaMarsh). The neural network is used to simulate a finite element analysis program to reduce the time and costs involved in obtaining an optimal design. The system, which executes on a Sun workstation, is written in Fortran and NETS is written in C and developed by NASA JSC.

DATA ACQUISITION SYSTEMS

(Norma K. Cambell, Charles H. Fox, Jr. and Lee T. McIntosh)

The data systems began in the mid-1960s at LaRC as dedicated on-site computer-based real time digital data acquisition, display, and control systems. The data system minicomputer was intended to have a full communication link to the central computer complex which treated data system processing support as second in priority to the real time flight simulators. Given the system limitations of the mid 1960s and the large number of facilities which were projected to acquire data systems, a decision was made not to support data systems from the central computer complex. Therefore, the data system computer took on the task of real time data reduction. The result was a monolithic data system design implemented on process control computers.

LaRC has a large number of diverse facilities which have data acquisition requirements. Throughout the 1970s and the 1980s, these facilities were equipped with data systems whose specifications still reflected the basic mid-1960s design concepts. Traditionally, researchers at a facility would define a set of requirements for the data acquisition needs of a facility. Based on these requirements, a data acquisition system was defined and developed for that facility by either the researchers themselves or by a group of data system specialists. These systems were one-of-a-kind systems. The data system was usually based on proprietary hardware and software consisting mainly of custom developed hardware interfaces and application software. From manpower considerations, it was highly desirable to minimize the number of unique proprietary hardware and software configurations. Facilities which had similar sets of requirements were therefore encouraged to choose a data system configuration that was already operational at another facility. This led to a situation in which there were several more or less standard data system configurations. In many cases, this new data system was obtained by sole source procurements based on the existence of the custom hardware interfaces and custom applications software.

There are two broad classes of facilities at LaRC: production facilities and research facilities. This distinction is not based on the size or complexity of either the facility or the data system. Production facilities are generally schedule driven and are in heavy use by NASA, industry, and other government organizations. Production facility data systems must be able to quickly and reliably acquire the data to meet such commitments. There is usually a high degree of standardization of hardware and software at the production facilities. In contrast, research facilities are not schedule driven and must be extremely flexible in order to be able to acquire data from non-standard and one-of-a-kind instrumentation. The development of new techniques and new instrumentation usually occurs in the research facilities. As these techniques and instrumentation are perfected, they migrate to the production facilities.

Current Data Systems

Figure 1 shows a schematic of the current system configuration for NTF. This system is based on a very tightly coupled proprietary

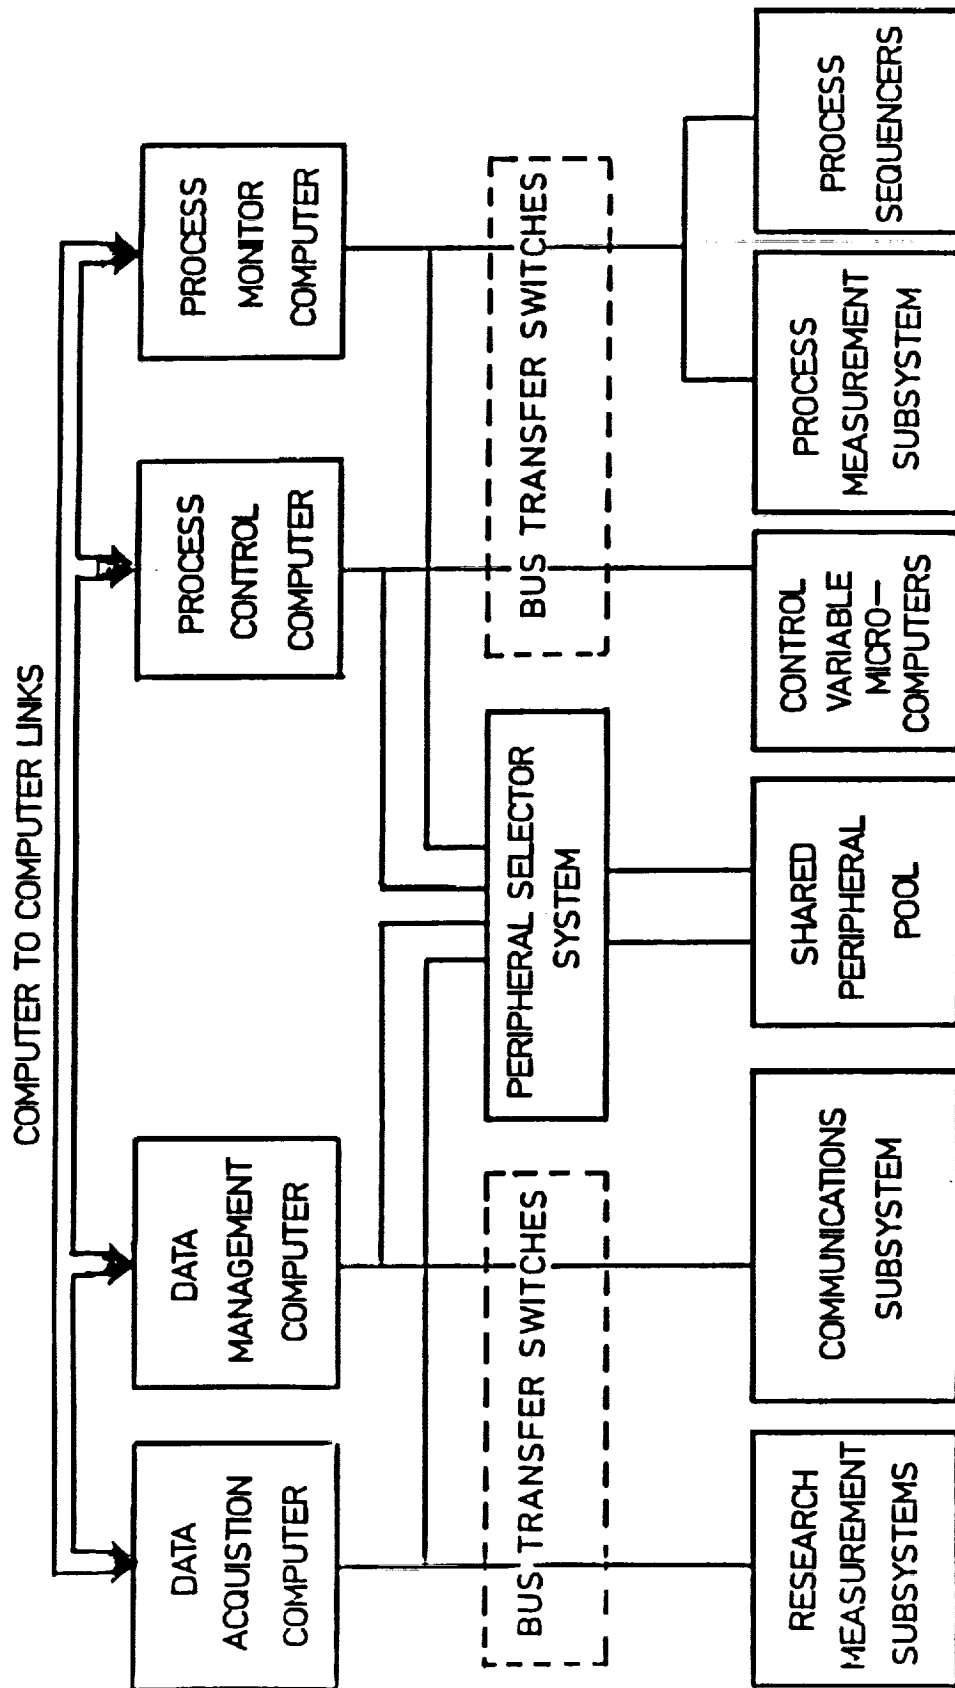


Figure 1. Current system configuration for NTF.

network of four 16-bit mini-computers. It is an older data system which is representative of the monolithic design concept employed during the 1970s and 1980s. More recent data systems use a 32-bit computer in the same basic design.

This type of proprietary data system configuration was not always fully satisfactory. In addition, the data system update or replacement cycle was approximately eight years. Being budget driven, computer technology changed very rapidly and systems were often obsolete prior to the next scheduled update.

In recent years, the analytical research computers have shifted from proprietary operating systems to UNIX operating systems, with many researchers having access to UNIX workstations. Data systems are migrating to open architectures and to UNIX systems with real time extensions.

The dynamically expanding capabilities of computer systems in recent years coupled with the emergence of open architecture designs has led to a fundamental reassessment of data system design. The efforts by several groups at LaRC have resulted in a new functional definition of what constitutes a research data system. These groups were composed of both data system providers and data system users who worked together. The final functional definition is therefore generic and applicable to all facilities and will form the basis of the data systems of the 1990s and beyond (figure 2).

The movement to open architecture was driven mainly by technical advances in the computer market as well as some of the following reasons: 1) procurement procedures forcing the issue to foster competition, 2) software portability, 3) hardware connectivity, 4) multi-processor boards and 5) third party software packages. Another significant driving force in this functional definition was economic. The change to a data system maintenance charge back to the facility in which the maintenance charge is a flat percent of the original capital investment made every facility extremely sensitive to the life cycle cost of the data systems.

Future Data Systems

As shown in figure 2, every new data system will include the following functions: data acquisition, data conversion, facility and

Data System: Organizational Schematic

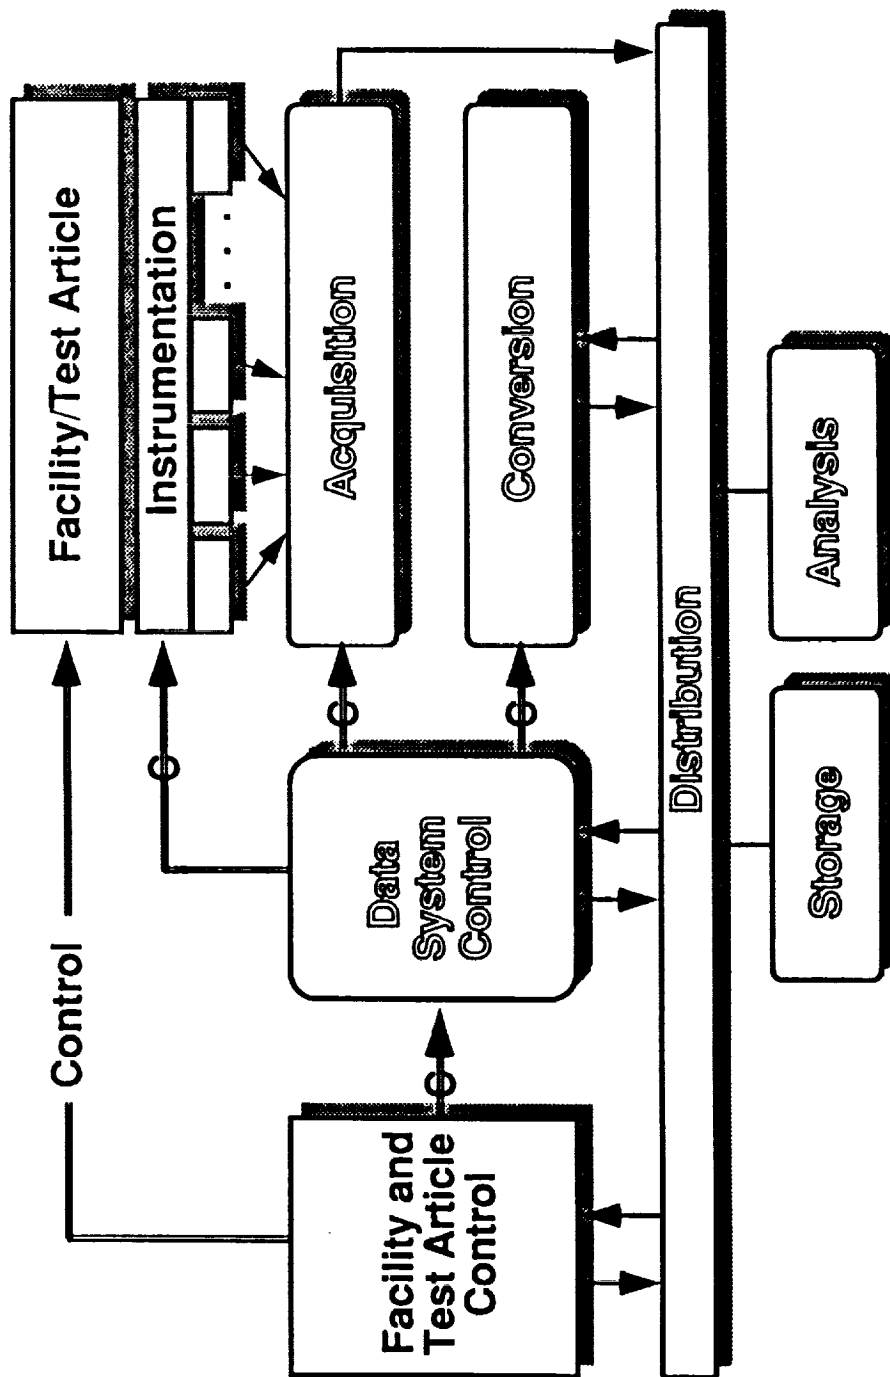


Figure 2. Future generic functional system configuration.

test article control, data system control, data distribution, data analysis, and data storage. Rather than having a monolithic hardware and software system to perform all of these functions, this concept envisions a distributed network of specialized hardware and modular software performing these separate functions. This new definition will accommodate the changes in computer systems and the tremendous advances in instrumentation. Future data acquisition systems will require interfaces to advanced instrumentation for the collection and analysis of fluid dynamic data. Advanced testing techniques will require large quantities of data as researchers move off the surface of the models and start measuring the fluid flow field surrounding test models. Advanced technologies will require faster computational speeds, and more mass storage. Conventional testing techniques will continue to be used. Third party graphical display packages will play a major role in the analysis of such large quantities of data as well as operator interfaces to the data acquisition systems. New analysis techniques will be required in the processing of large quantities of data from thermographic, phosphor and thin film types of data. For instance, just one pixel on just one frame of a thermographic image can represent a single data point. Thus one can easily see the massiveness of data which will be required to be saved and/or processed in the future.

General requirements of the future data system are that the computer hardware and instrumentation will be open architecture, scaleable, reliable, and easy to maintain. Both hardware and software standards will be used as much as possible. The research data system will satisfy near-real time requirements although it may contain embedded modular real time subsystems to satisfy a time critical requirement. The data system shall incorporate new technology through simple upgrades rather than total replacement. The software environment is expected to reduce the workload on the researcher through software libraries, extensive use of graphical user interfaces, easily re-configurable display screens, and standard interfaces. The data system will be able to be readily configurable to satisfy security requirements for classified data acquisition and processing when necessary.

The functional concept and open architecture applied to a proposed data system for NTF is illustrated in figure 3. This concept represents a loosely coupled collection of computers and workstations that each perform assigned functions and communicate

over a local network. Advancements made in communication will play a key roll in the operation of future data acquisition systems.

Problem Areas

The researcher and technician do not use just one facility. They routinely use different facilities. Therefore, there is a requirement for the data system at each facility to present a similar user interface to both the researcher and the technician.

Procurement regulations have not kept up with the rapidly changing computer and instrument technology advancements. The procurement cycle is often longer than the technology update cycle. Procurement regulations do not permit a series of separate procurements to obtain the same standard data system configuration; in fact, for a series of open architecture purchases, procurement regulations would almost guarantee that there was no commonality in the data systems. This poses a nightmare for training, maintenance, and configuration control.

The future data systems will consist of a cooperating network of computers running a variety of software packages from both the manufacturers and from numerous third party sources. To individually license the software packages on each of these computers could become extremely expensive. For this reason, LaRC must be able to negotiate site license contracts for each of these software packages which treat LaRC as one site with appropriate volume discounts. Another software problem area is that embedded microprocessors are frequently used in equipment and vendors rarely provide the source code running on these microprocessors.

There is another potential problem in open architecture procurements from a specification viewpoint. You can obtain the same item from different vendors, but they may not be electrically or functionally compatible or interchangeable.

For major facilities, the update cycle is dependant on the IMS data system update budget, which has in the past funded one or two systems each year over an eight year cycle. The major facilities list includes almost all of the production facilities and some research facilities. The budgeting process makes insufficient provisions for upgrading data systems in research facilities and has

yet to address a massive technology-driven upgrade for all data systems at the same time.

The development and installation of new open architecture data systems, while maintaining the old proprietary data systems, will require additional manpower resources. Likewise a massive upgrade of all data systems in a timely fashion would require additional personnel. However, the new data systems would significantly contribute to the research productivity of the facilities. Thus it would be desirable to have a productivity line item in the budget that would allow an indefinite quantity procurement of new data systems for the major facilities.

VISUALIZATION

**(John T. Bowen, C. Mark Cagle, Clyde R. Gumbert,
Kurt Severance, and Kathryn Stacy)**

Scientific visualization plays an important role in many diverse research disciplines at LaRC, including structures, computational fluid dynamics (CFD), experimental flow field measurements, remote sensing and atmospheric sciences, and robotics. Each of these application areas have in turn played a part in driving the current state of the art of visualization at LaRC. This section will address the current computing systems and software tools available for visualizing scientific data. Several important issues in the field of visualization will also be discussed. Finally, potential areas of advancement of LaRC's visualization capabilities will be discussed in light of future trends in the Center's research disciplines and in the visualization field itself.

Data Visualization and Animation Lab

Several research facilities have their own local visualization capabilities. However, since many visualization techniques require specialized hardware, training and experience, few individual research facilities can afford their own. The primary facility at LaRC for visualizing scientific data is the Data Visualization and Animation Lab (DVAL) in ACD. DVAL consists of key personnel, advanced software packages, and state-of-the-art hardware for investigating and evaluating scientific data. Capabilities of the lab include digital image processing, computer graphics, data visualization and analysis, solid modeling, and video editing and

special effects for effective dissemination of results. The most effective use of these capabilities is via the interdisciplinary team approach, where the engineer/scientist works together with the visualization expert to determine and perform the most effective means of graphically extracting pertinent information from the data.

Computational Fluid Dynamics

Computational fluid dynamics codes have been developed to study how changes in the geometry and flow variables affect flow fields. Interpretation of CFD results requires advanced visualization tools. A high-end Silicon Graphics, Inc. (SGI) Iris graphics workstation is the key platform used throughout the Center for visualizing this data. The Flow Analysis Software Toolkit (FAST), developed by Sterling Software under contract with NASA Ames, is the primary environment on the IRIS for viewing and interpreting CFD results. This toolkit is capable of displaying contours, isosurfaces, particle traces, vector fields, and scalar function surfaces. It also provides tools for rendering complete animations with titles and legends. A sample is shown in figure 4. Similar capabilities are provided by several commercially available packages such as Fieldview, AVS, TECPLOT, and PV~WAVE. While these systems are useful for conveying overall information about one's data, the capability of probing and extracting quantitative information is often lacking.

Flow Visualization Techniques

Another source of data at LaRC is from flow visualization techniques performed at the various experimental facilities. Flow visualization techniques provide qualitative information about the flow field under observation. Commonly, a video camera is positioned outside of wind tunnel windows, and the experiment is recorded onto videotape. DVAL supports specialized hardware and software referred to as the Video Image Processing System (VIPS) for post-processing such videotapes. The system is designed to digitize, process, store, and retrieve frames at real-time video rates, or 30 frames per second. Special features of the VIPS system include the ability to perform frame accurate digitization based on timecode, the ability to store several minutes of digital data, and the ability to perform image processing operations such as edge detection, spatial convolution, and arithmetic and logic operations at video rates. Recent techniques have also been developed to extract 3D flow field information from the videotapes given

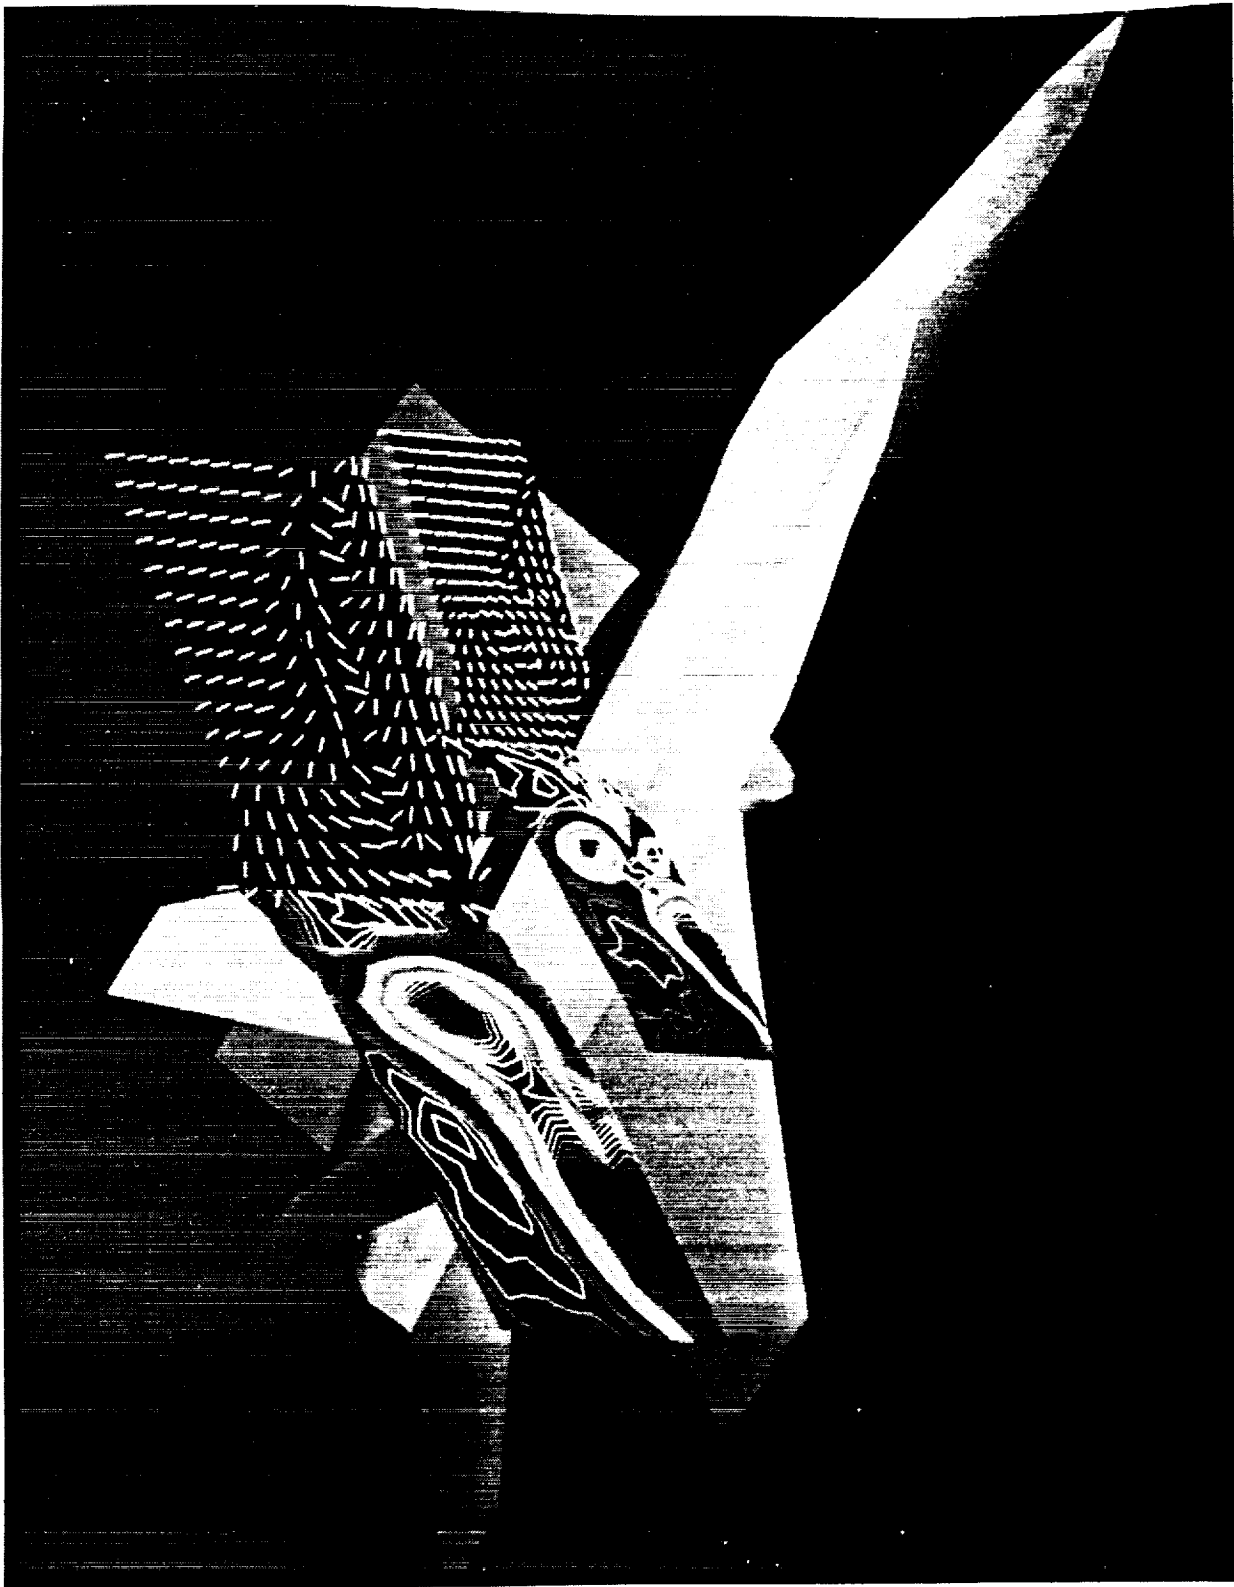


Figure 4. Visualization of experimental laser velocimetry data using the Flow Analysis Software Toolkit.

knowledge of the camera and lighting systems used during the experiment.

Rendering and Animation

Researchers also require photo-realistic rendering and advanced animation tools to turn raw data, such as computed parameters describing the rotation of space station solar panels, into broadcast quality videos for dissemination. The main package used by DVAL to create these animations is WAVEFRONT's Advanced Visualizer which runs on an Iris workstation. A sample is shown in figure 5. WAVEFRONT contains a solid modeler to build the objects in the animation, and can be used, for example, to model the components of a space station. Once the model has been defined, textures, lighting, and surface attributes are chosen to create realistic scenes. Finally, keyframes in the animation are specified to create a series of images depicting smooth motion. This system has been used at LaRC to simulate satellite retrievals, robot arm deflections, remote-sensing methods, and wind tunnel configurations.

Video Editing and Dissemination

A broadcast quality video editing and special effects system is necessary to provide an effective means for disseminating information. Animations of dynamic processes produced by the WAVEFRONT package, for example, can be recorded to video along with title pages and audio tracks. Fully digital editing systems consisting of real-time disks, digital video recorders, and special effects hardware are now available at LaRC. This type of system allows images to be transmitted over existing networks, and allows multiple editing sessions without degradation of the original video frames. Essentially anything that can be displayed on the computer screen can now be recorded onto videotape, allowing results to reach a wider audience. The effectiveness of technical video dissemination is highly recognized at LaRC, and efforts are currently underway to develop standards for a NASA video publication.

Comparison of Experimental and Computational Data

An ongoing area of study is the effective comparison of computational and experimental results within the same visualization environment. Unfortunately, few visualization packages directly support the inherent differences between these

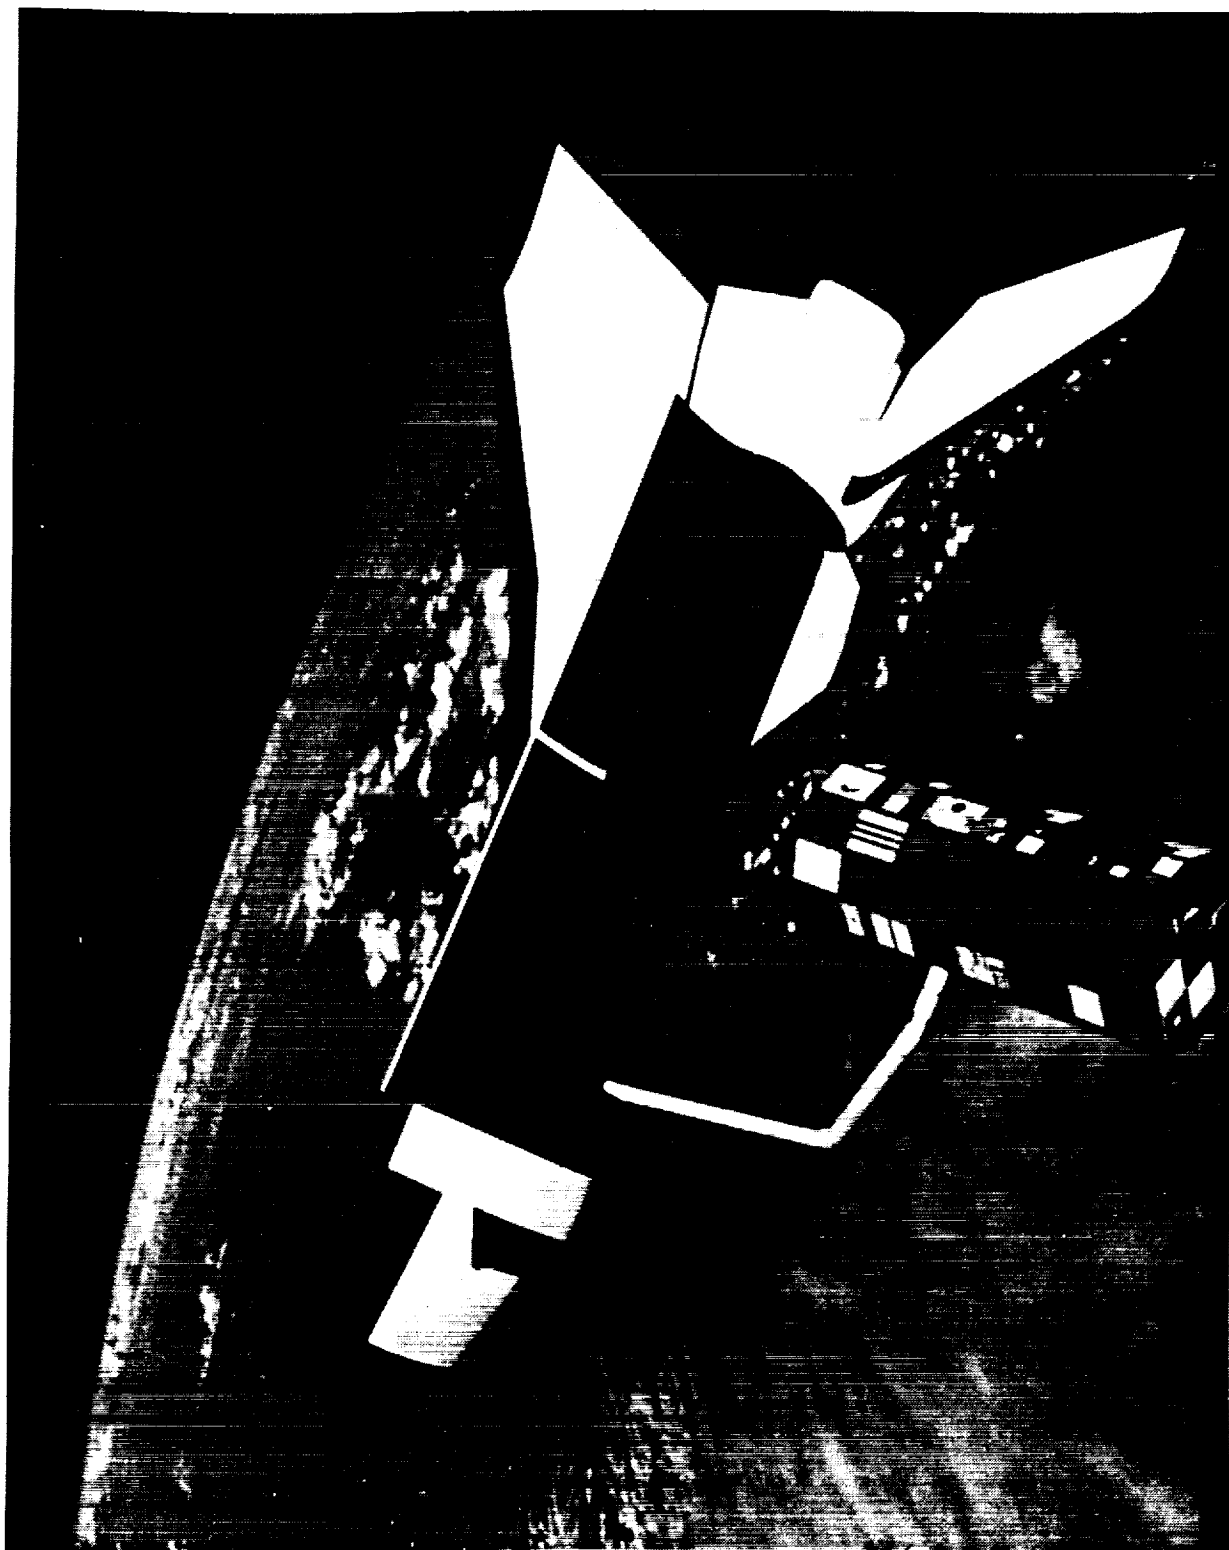


Figure 5. Long Duration Exposure Facility retrieval simulation rendered by the Wavefront Advanced Visualizer.

two forms of data. Modern CFD techniques, for example, can generate the pressure, temperature, or density at any point in the flowfield. Wind tunnel tests, on the other hand, generally provide data in the form of video images which record the intensities of backscatter arriving at a camera from an illuminated flow field. Few tools are available to quantitatively compare the CFD data with this image intensity information.

Several promising experimental flow visualization techniques have recently been developed, though, which facilitate the quantitative validation of CFD codes. Researchers now have the capability of measuring more than just forces, moments, and limited surface phenomena. Laser induced fluorescence, for example, makes use of fluorescent tracer materials and has the potential of providing density, temperature, and velocity for a cross-section of the flow. In addition, doppler velocimetry methods have proven useful for obtaining velocity data. DVAL is actively involved in developing tools to facilitate the comparison of these experimental results with those from CFD.

Standards

Graphics standards affect how images are composed and displayed on a computer screen and how graphics metafiles are generated and postprocessed for hardcopy output. Early graphics standards defined by ANSI such as CORE, GKS, and PHIGS have led to traditional graphics libraries like Common Graphics Library, DI-3000, and NCAR which are still heavily used for two- and three-dimensional graphics display. These graphics subroutine libraries are steadily being replaced by packages like TECPLOT, PV~WAVE, FIELDVIEW, and AVS which are based on the X-Windows standard and are portable across workstation platforms. The common denominator among these new visualization programs is the ability to provide network-transparent, point-and-click, visual data analysis with an intuitive graphical user interface (GUI). It is anticipated that the Motif and X-windows standards will drive the development of graphics software in the near future.

Standards are also needed to efficiently exchange geometric information between the design and analysis phases of engineering. CAD/CAM vendors, for example, use the IGES data standard to transfer information among the CAD programs, whereas CFD packages generally use an incompatible format such as the PLOT3D

standard to describe complex geometries and the computational volumes around them. An effort is currently under way at NASA LaRC, Lewis, and Ames to develop a subset of the IGES standard to permit efficient communication of surface data to and from CAD and CFD programs. The NASA standard is focusing on NURBS (Non-Uniform Rational B Splines) surfaces as the surface representation of choice for aircraft. Routines are currently being obtained to convert any surface representation from CAD systems to NURBS surfaces for use in CFD codes. Unfortunately, most of the current CFD codes will need to be modified to accommodate NURBS surfaces.

Applications Driving Future Needs

Several applications at LaRC will demand more advanced visualization tools in the near future. New CFD codes, for example, are being developed to examine chemical composition and turbulence within supersonic flows, as in scramjet engines. Techniques are not yet available to effectively analyze these parameters. In the area of atmospheric science, significant efforts will be necessary to compress, organize, and visualize the terabytes of data transmitted from the EOS platform on a daily basis. In addition, the wind tunnels are being configured with devices, such as doppler global velocimeters, which extract 3D quantitative flow field data for which no standard visualization tools are available. Finally, the extraction of 3D information from lengthy video recordings of flow field experiments requires extensive processing power, storage space, and rendering capabilities which are not yet available.

Future Trends in Visualization Technology

Advances in the technology of visualization will have an effect on how information is graphically extracted from data in terms of the methods in which it is extracted and the mechanisms through which it is perceived.

There is an inherent loss of information when a three dimensional scene is displayed on a two dimensional surface such as a computer monitor screen or a piece of paper. Simple techniques to replace the illusion of depth, such as perspective transformations and intensity mappings as a function of depth have been used for a long time. Holographic techniques are being developed to create images which appear to be three dimensional. However these techniques are still too expensive in terms of both time and money to be used routinely

by the researcher. Stereo viewing is a technique which has been tried in several forms over the last several years, primarily in the cinema, with varying degrees of success. As graphics workstation performance increases, the ability to display an image for each eye at a rate sufficient to create the illusion of three dimensionality without a great deal of eye-strain improves. Beyond stereo viewing techniques is "Virtual Reality" where stereo is a means of display which gives an improved perception of a three dimensional scene. Virtual reality is a user interface technique which uses many mechanisms to simulate the acts of looking around and interacting with a three dimensional environment. "Virtual reality...requires an earnest focus on the user interface of the application. Designers must consider the human factors of the devices which sense and stimulate the experience, the tasks to be performed, and the desired experience. Virtual reality applications enable synthetic visualizations and experiences which may involve real data spaces and real people, or may present entirely simulated environments. In order to effectively encourage creativity and productivity, the virtual experience must be credible." (SIGGRAPH '91)

Recent advances in networking and Graphical User Interfaces (GUI's) have opened up new areas for improvements in visualization techniques. Visual programming systems such as AVS, Khoros and SGI Explorer make use of inter-process communication protocols and graphics to allow the user to connect together several individual modules that each perform a specific function in order to develop a comprehensive application. As inter-computer communications networks become more 'transparent' to the user, these individual tasks can be performed on the computer best suited to that particular module.

As visualization tools become increasingly interactive and multi-dimensional, more effective means for disseminating the scientific results are being sought. The hardware to convert sequences of computer-generated images into standard NTSC video have long been available, but the resulting animations often lose the fidelity of the original images and require viewers to scan information sequentially. Emerging multimedia PC platforms, however, allow users to navigate at their own pace through an integrated environment of text, graphics, video, sound, and raw data. This information can be widely distributed on an inexpensive CD-ROM which provides an enormous storage capacity of 600MB. What was

once limited to a few public libraries is now widely available even for home use in the form of interactive CD-ROM magazines, encyclopedias, documentaries, literary works, and scientific databases. The aerospace research community is just beginning to investigate the feasibility of this new technology, and LaRC is currently pursuing a prototype for multimedia-based technical reports.

Before CD-ROM technology becomes a viable medium for distributing scientific results, several limitations must be addressed, including publication costs, access speed, and standardization. While the price of individual CD-ROM drives and disks continues to decline, the expertise and expense required to produce the original master can be staggering. In addition, the relatively slow access speed of a standard CD-ROM (150 Kbytes/sec) prevents the effective playback of large video sequences. More effective video compression ratios and decompression rates are necessary for CD-ROM drives to deliver the necessary performance. Both Apple and the Motion Picture Experts Group are pursuing such techniques. Finally, hardware and software vendors must agree on a standard format before CD-ROM's can be widely distributed. Sony, Philips, Intel, and Microsoft are the leading corporations which will help define this standard.

SUMMARY

Each section presents details about what the future holds for their area. In general, there is a need for real-time systems to support piloted real time simulation, embedded flight software, and data acquisition. Researchers requiring high performance computing should embrace multiple (possibly massively parallel) CPU's to achieve the necessary computational speed for solving their type of problems. Other areas such as embedded systems, artificial intelligence, and visualization not only must examine the computing power, but also examine standards, quality development tools, languages, operating systems, and portability. Common threads running through these areas include standards such as C or ADA for languages, UNIX for operating systems, and X-windows. Adopting these standards should also aid in portability. The problems with data storage and transfer become apparent in the sections on visualization and data acquisition systems. Almost all areas mention the need to communicate among different computers. Researchers at their desk need to easily access an appropriate

computing environment. These problems can only be solved if the proper distributed computing network is available. Thus for LaRC to remain a state-of-the art computing center, we must keep advancing as we have in the past and take advantage of as many of the new hardware and software developments as the budget will allow.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 1992	3. REPORT TYPE AND DATES COVERED Technical Memorandum	
4. TITLE AND SUBTITLE Current State and Future Direction of Computer Systems at NASA Langley Research Center			5. FUNDING NUMBERS WU 509-10-11-04	
6. AUTHOR(S) Computer Systems Technical Committee				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NASA Langley Research Center Hampton, VA 23681-0001			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001			10. SPONSORING / MONITORING AGENCY REPORT NUMBER NASA TM-107685	
11. SUPPLEMENTARY NOTES This report was edited by James L. Rogers and Jerry H. Tucker of Langley Research Center.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Unclassified - Unlimited Subject Category 61			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Computer systems have advanced at a rate unmatched by any other area of technology. As performance has dramatically increased there has been a no less dramatic reduction in cost. This constant cost performance improvement has precipitated the pervasiveness of computer systems into virtually all areas of technology. This improvement is due primarily to advances in microelectronics. Most people are now convinced that the new generation of super computers will be built using a large number (possibly thousands) of high performance microprocessors. Although the spectacular improvements in computer systems have come about because of these hardware advances, there has also been a steady improvement in software techniques. In an effort to understand how these hardware and software advances will effect research at NASA Langley Research Center, the Computer Systems Technical Committee drafted this white paper to examine the current state and possible future directions of computer systems at the Center. This paper discusses selected important areas of computer systems including real-time systems, embedded systems, high performance computing, distributed computing networks, data acquisition systems, artificial intelligence, and visualization.				
14. SUBJECT TERMS real time computer systems, embedded systems, high performance computing, distributed computing networks, artificial intelligence, data acquisition systems, visualization			15. NUMBER OF PAGES 33	
			16. PRICE CODE A03	
17. SECURITY CLASSIFICATION OF REPORT unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT unclassified	20. LIMITATION OF ABSTRACT	